Fast separation in a graph with an excluded minor †

Bruce Reed[1] and David R. Wood[2]

[1]*School of Computer Science, McGill University, Montréal, Canada* (breed@cs.mcgill.ca)
[2]*Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Spain* (david.wood@upc.edu)

Let G be an n-vertex m-edge graph with weighted vertices. A pair of vertex sets $A, B \subseteq V(G)$ is a $\frac{2}{3}$-*separation* of *order* $|A \cap B|$ if $A \cup B = V(G)$, there is no edge between $A \setminus B$ and $B \setminus A$, and both $A \setminus B$ and $B \setminus A$ have weight at most $\frac{2}{3}$ the total weight of G. Let $\ell \in \mathbb{Z}^+$ be fixed. Alon, Seymour and Thomas [*J. Amer. Math. Soc.* 1990] presented an algorithm that in $\mathcal{O}(n^{1/2}m)$ time, either outputs a K_ℓ-minor of G, or a separation of G of order $\mathcal{O}(n^{1/2})$. Whether there is a $\mathcal{O}(n + m)$ time algorithm for this theorem was left as open problem. In this paper, we obtain a $\mathcal{O}(n + m)$ time algorithm at the expense of $\mathcal{O}(n^{2/3})$ separator. Moreover, our algorithm exhibits a tradeoff between running time and the order of the separator. In particular, for any given $\epsilon \in [0, \frac{1}{2}]$, our algorithm either outputs a K_ℓ-minor of G, or a separation of G with order $\mathcal{O}(n^{(2-\epsilon)/3})$ in $\mathcal{O}(n^{1+\epsilon} + m)$ time.

Keywords: graph algorithm, separator, minor

1 Introduction

We consider graphs G that are simple, finite, and undirected. Let $V(G)$ and $E(G)$ denote the vertex and edge sets of G. Let $|G| := |V(G)|$ and $\|G\| := |E(G)|$. A *separation* of G is a pair $\{A, B\}$ of vertex sets $A, B \subseteq V(G)$ such that $A \cup B = V(G)$, and there is no edge with one endpoint in $A \setminus B$ and the other endpoint in $B \setminus A$. The *order* of $\{A, B\}$ is $|A \cap B|$. The set $A \cap B$ is called a *separator* of G. A *weighting* of G is a function $w : V(G) \to \mathbb{R}^+$. Let $w(S) := \sum_{v \in S} w(v)$ for all $S \subseteq V(G)$, and $w(G) := w(V(G))$. We say (G, w) is a *weighted graph*. A separation $\{A, B\}$ of a weighted graph (G, w) is an α-*separation* if $w(A \setminus B) \leq \alpha \cdot w(G)$ and $w(B \setminus A) \leq \alpha \cdot w(G)$.

A 'separator theorem' is of the format: for some $0 < \alpha < 1$ and $0 < \epsilon \leq 1$, every graph G from a certain family has an α-separation of order $\mathcal{O}(|G|^{1-\epsilon})$. Applications of separator theorems are numerous, and include VLSI circuit layout, approximation algorithms using the divide-and-conquer paradigm, solving sparse systems of linear equations, pebbling games, and graph drawing. See the recent monograph by Rosenberg and Heath [9] for more details.

A seminal theorem due to Lipton and Tarjan [5] states that every weighted planar graph G has a $\frac{2}{3}$-separation of order $\mathcal{O}(|G|^{1/2})$ that can be computed in $\mathcal{O}(|G| + \|G\|)$ time. This result was generalised

†Research of B.R. is supported by NSERC. Research of D.W. is supported by the Government of Spain grant MEC SB2003-0270.

for graphs with an excluded minor by Alon *et al.* [1] (see [2, 3, 7] for related results). A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges, in which case we say that G has an H*-minor*. The Kuratowski-Wagner Theorem states that a graph is planar if and only if it has no K_5-minor and no $K_{3,3}$-minor. An H*-model* in G is a set of disjoint connected subgraphs $\{X_v : v \in V(H)\}$ indexed by the vertices of H, such that for every edge $vw \in E(H)$, there is an edge $xy \in E(G)$ with $x \in X_v$ and $y \in X_w$. Clearly G has an H-minor if and only if G has an H-model. We choose to work with H-models rather than H-minors.

Theorem 1 (Alon *et al.* [1]) *There is an algorithm with running time $\mathcal{O}((\ell \cdot |G|)^{1/2} \cdot (|G| + \|G\|))$ that, given $\ell \in \mathbb{Z}^+$ and a weighted graph (G, w), either outputs:*
(a) *a K_ℓ-model of G, or*
(b) *a $\frac{2}{3}$-separation of (G, w) of order at most $\ell^{3/2} \cdot |G|^{1/2}$.*

Suppose that ℓ is fixed. It follows from a result of Mader [6] (see Theorem 3) that Theorem 1 can be implemented in $\mathcal{O}(|G|^{3/2} + \|G\|)$ time. Alon *et al.* [1] left as an open problem whether linear time is possible. The main result of this paper is the following partial answer to this question. We obtain a linear running time at the expense of a slightly larger separator (and a larger dependence on ℓ). Moreover, our algorithm exhibits a tradeoff between running time (ranging from $\mathcal{O}(n)$ to $\mathcal{O}(n^{3/2})$) and the order of the separator (ranging from $\mathcal{O}(n^{2/3})$ to $\mathcal{O}(n^{1/2})$).

Theorem 2 *There is an algorithm with running time $\mathcal{O}(2^{(3\ell^2+7\ell-3)/2} \cdot |G|^{1+\epsilon} + \ell \cdot \|G\|)$ that, given $\epsilon \in [0, \frac{1}{2}]$, $\ell \in \mathbb{Z}^+$, and a weighted graph (G, w), either outputs:*
(a) *a K_ℓ-model of G, or*
(b) *a $\frac{2}{3}$-separation of (G, w) of order at most $2^{(\ell^2+3\ell+1)/2} \cdot |G|^{(2-\epsilon)/3}$.*

Note that for applications to divide-and-conquer algorithms a separation of order $\mathcal{O}(|G|^{1-\epsilon})$, for some constant $\epsilon > 0$, is all that is needed.

The idea behind the proof of Theorem 2 is simple. We now outline the proof for fixed ℓ and with $\epsilon = 0$. Suppose that in $\mathcal{O}(|G| + \|G\|)$ time, we can find a partition of $V(G)$ into $|G|^{2/3}$ connected subgraphs $\{S_1, S_2, \dots, S_{|G|^{2/3}}\}$, each containing $\mathcal{O}(|G|^{1/3})$ vertices. Let H be the weighted graph obtained from G by contracting each S_i to a vertex v_i with weight $w(v_i) = w(S_i)$. Then apply Theorem 1 to H to either obtain a K_ℓ-model in H which defines a K_ℓ-model in G, or a $\frac{2}{3}$-separation $\{A, B\}$ of H with order $\mathcal{O}(|H|^{1/2}) = \mathcal{O}(|G|^{1/3})$, in which case $\{\bigcup\{S_i : v_i \in A\}, \bigcup\{S_i : v_i \in B\}\}$ a is a $\frac{2}{3}$-separation of G with order $\mathcal{O}(|G|^{2/3})$. The running time is $\mathcal{O}(|H|^{3/2} + \|H\|) \subseteq \mathcal{O}(|G| + \|G\|)$. The proof of Theorem 2 is actually a little different from this outline. In particular, the subgraphs S_i will not necessarily be connected, but we will still be able to convert the output from Theorem 1 applied to H to the desired output for G. By relaxing the connectivity condition, we are able to prove that an appropriate partition exists.

We will use the following notation for a graph G. For $x \in V(G)$, let $N(x) := \{y \in V(G) : xy \in E(G)\}$. For a subgraph X of G, let $N(X) := \bigcup\{N(x) \setminus X : x \in X\}$. Where there is no confusion, a set of vertices $S \subseteq V(G)$ will also refer to the subgraph of G induced by S.

2 Mader's Theorem

This section contains a number of easily proved results—see the full version of the paper for details. We start with an algorithmic version of a theorem of Mader [6] (cf. [8]).

Theorem 3 *Given a graph G with $\|G\| \geq 2^{\ell-3} \cdot |G|$ (for some $\ell \in \mathbb{Z}^+$), a K_ℓ-model of G can be computed in $\mathcal{O}(\ell(|G| + \|G\|))$ time.*

Note that if we ignore the running time, Theorem 3 is far from best possible. Kostochka [4] and Thomason [10] independently proved that if $\|G\| \in \Omega(\ell\sqrt{\log \ell} \cdot |G|)$ then G has a K_ℓ-model. Theorem 3 implies the following slightly faster version of Theorem 1 (for fixed ℓ)

Theorem 4 *There is an algorithm with running time $\mathcal{O}(2^{2\ell} \cdot |G|^{3/2} + \ell \cdot \|G\|)$ that, given $\ell \in \mathbb{Z}^+$ and a weighted graph (G, w), either outputs:*
(a) *a K_ℓ-model of G, or*
(b) *a $\frac{2}{3}$-separation of (G, w) of order at most $\ell^{3/2} \cdot |G|^{1/2}$.*

A *k-clique* of G is a (not necessarily maximal) set of k pairwise adjacent vertices of G. If every subgraph of G has a vertex of degree at most d, then G is *d-degenerate*. For example, Theorem 3 implies that a graph with no K_ℓ-minor is $2^{\ell-2}$-degenerate. It is easily proved that a d-degenerate graph G with no k-clique has less than $d^{k-1} \cdot |G|$ cliques. Hence a graph G with no K_ℓ-minor has less than $2^{(\ell-2)(\ell-1)} \cdot |G|$ cliques. For an algorithm, we have the following result.

Lemma 1 *Given a graph G with no k-clique and at least $2^{(\ell-2)(k-1)} \cdot |G|$ cliques (for some $\ell \in \mathbb{Z}^+$), a K_ℓ-minor of G can be computed in $\mathcal{O}(\ell(|G| + \|G\|))$ time.*

3 Proof of Theorem 2

Let G be a graph. Let $\mathcal{A}$ be a partition of $V(G)$. Let H be the graph obtained from G by collapsing each part $S \in \mathcal{A}$ to a single vertex v, and replacing parallel edges by a single edge. Denote $H_v := S$. We say $\{H_v : v \in V(H)\}$ is an *H-partition* of G. Furthermore, $\{H_v : v \in V(H)\}$ is a *connected* H-partition of G if $vw \in E(H)$ if and only if there is an edge of G between every component of H_v and every component of H_w. We prove the following lemma.

Lemma 2 *There is an algorithm with running time $\mathcal{O}(2^{2\ell} \cdot |G| + \|G\|)$ that, given $\ell, k \in \mathbb{Z}^+$ and a graph G, outputs a connected H-partition of G such that either:*
(a) *H has a K_ℓ-model (which is also output), or*
(b) *$|H| \leq 2^{\ell^2+\ell-1} \cdot |G| \cdot k^{-1}$, and $|H_x| < 2k$ for all $x \in V(H)$.*

Proof of Theorem 2 assuming Lemma 2: Apply Lemma 2 with $k = \lfloor |G|^{(1-2\epsilon)/3} \rfloor$. First suppose that Lemma 2 outputs a K_ℓ-model $\{S_1, S_2, \ldots, S_\ell\}$ of H. Thus each S_i is a connected subgraph of H. Choose a connected component Z_v of H_v for each $v \in V(H)$. Let $T_i := \bigcup\{Z_v : v \in S_i\}$. Then $\{T_1, T_2, \ldots, T_\ell\}$ is a K_ℓ-model of G.

Otherwise $|H| \leq 2^{\ell^2+\ell-1} \cdot |G|^{2(1+\epsilon)/3}$, and $|H_x| < 2|G|^{(1-2\epsilon)/3}$ for all $x \in V(H)$. Let $w(v) := w(H_v)$ for all $v \in V(H)$. Apply Theorem 4 to (H, w). The running time is

$$\mathcal{O}(2^{2\ell} \cdot |H|^{3/2} + \ell \cdot \|H\|) \subseteq \mathcal{O}(2^{2\ell} \cdot (2^{\ell^2+\ell-1} \cdot |G|^{2(1+\epsilon)/3})^{3/2} + \ell \cdot \|G\|) \subseteq \mathcal{O}(2^{(3\ell^2+7\ell-3)/2} \cdot |G|^{1+\epsilon} + \ell \cdot \|G\|) \ .$$

We either obtain a K_ℓ-model of H, or a $\frac{2}{3}$-separation of H with order at most $\ell^{3/2} \cdot |H|^{1/2}$. In the first case, G has a K_ℓ-model as proved above.

Now suppose that we obtain a $\frac{2}{3}$-separation $\{A, B\}$ of (H, w) with order

$$|A \cap B| \leq \ell^{3/2} \cdot |H|^{1/2} \leq \ell^{3/2} \cdot (2^{\ell^2+\ell-1}|G|^{2(1+\epsilon)/3})^{1/2} \leq 2^{(\ell^2+3\ell-1)/2} \cdot |G|^{(1+\epsilon)/3} \ .$$

Let $X := \bigcup\{H_v : v \in A\}$ and $Y := \bigcup\{H_v : v \in B\}$. Then $\{X, Y\}$ is a separation of G with order

$$|X \cap Y| = |\bigcup\{H_v : v \in A \cap B\}| \leq 2^{(\ell^2+3\ell-1)/2} \cdot |G|^{(1+\epsilon)/3} \cdot 2|G|^{(1-2\epsilon)/3} \leq 2^{(\ell^2+3\ell+1)/2} \cdot |G|^{(2-\epsilon)/3} \ .$$

We have $w(X \setminus Y) = w(A \setminus B) \leq \frac{2}{3}w(H) = \frac{2}{3}w(G)$. Similarly $w(B \setminus A) \leq \frac{2}{3}w(G)$. □

Proof of Lemma 2:

Step 1: Using a breadth-first search algorithm, compute a maximal set $\mathcal{A}$ of connected subgraphs of G such that $|S| = k$ for all $S \in \mathcal{A}$. Let $\mathcal{B}$ be the set of connected components of $G \setminus \bigcup\{S \in \mathcal{A}\}$. Then $\mathcal{A} \cup \mathcal{B}$ is a partition of $V(G)$, and there is no edge of G between distinct $T_1, T_2 \in \mathcal{B}$. Note that $|T| < k$ for all $T \in \mathcal{B}$, as otherwise T would contain a connected subgraph X with $|X| = k$, which could be added to $\mathcal{A}$.

Step 2: Let H be the graph obtained from G by contracting each set $S \in \mathcal{A} \cup \mathcal{B}$ into a single vertex v with $H_v := S$, and replacing parallel edges by a single edge. Since each $S \in \mathcal{A} \cup \mathcal{B}$ is connected in G, $\{H_v : v \in V(H)\}$ is a connected H-partition of G. Let $A := \{v \in V(H) : H_v \in \mathcal{A}\}$ and $B := \{v \in V(H) : H_v \in \mathcal{B}\}$. A vertex v of H is *big* if $|H_v| \geq k$. A vertex v of H is *small* if $|H_v| < k$. Observe that every vertex in A is big, B is an independent set of H, and every vertex in B is small. Partition $B = C \cup D \cup E$, where $C := \{v \in B : \deg_H(v) \geq 2^{\ell-2}\}$, $D := \{v \in B : \ell - 1 \leq \deg_H(v) < 2^{\ell-2}\}$, and $E := \{v \in B : \deg_H(v) \leq \ell - 2\}$.

Step 3: Suppose that $|C| \geq |A|$. Then the subgraph $C \cup A$ of H has at least $2^{\ell-2} \cdot |C|$ edges and at most $2|C|$ vertices. By Theorem 3, a K_ℓ-model of $C \cup A$ can be computed in $\mathcal{O}(\ell \cdot |G|)$ time. We now assume that $|C| < |A|$.

Step 4: For each vertex $v \in D \cup E$, if there is a pair $x, y \in A$ of distinct neighbours of v, such that $\{x, y\}$ has not been assigned any vertex in $D \cup E$, then assign v to $\{x, y\}$. This step can be implemented in $\mathcal{O}(2^{2\ell} \cdot |G|)$ time, since each vertex in $D \cup E$ has degree at most $2^{\ell-2}$.

Suppose that there is a vertex $v \in D$ that is not assigned. Let the neighbourhood of v be $\{x_1, x_2, \dots, x_d\}$. Then $d \geq \ell - 1$. Thus for all $1 \leq i < j \leq d$, there is a distinct vertex $v_{i,j}$ that is assigned to the pair $\{x_i, x_j\}$, and $v_{i,j}$ is adjacent to both x_i and x_j. In the graph obtained from H by contracting each edge $x_i v_{i,j}$, the subgraph $\{x_1, x_2, \dots, x_d, v\}$ is a clique on at least ℓ vertices. Thus H has a K_ℓ-model. We now assume that every vertex in D is assigned.

Let E^* be the set of assigned vertices in E. Consider the graph obtained from $A \cup D \cup E^*$ by contracting the edge vx for each $v \in D \cup E^*$ assigned to the pair $\{x, y\}$. This graph has $|A|$ vertices and at least $|D| + |E^*|$ edges. Thus if $|D| + |E^*| \geq 2^{\ell-3} \cdot |A|$, then by Theorem 3, H has a K_ℓ-model that can be computed in $\mathcal{O}(\ell \cdot |G|)$ time. We now assume that $|D| + |E^*| < 2^{\ell-3} \cdot |A|$.

Step 5: Partition $E \setminus E^* = \bigcup\{P_1, P_2, \dots, P_s\}$ such that for all $u, v \in E \setminus E^*$, we have $N(u) = N(v)$ if and only if both $u, v \in P_i$ for some $1 \leq i \leq s$. For all $1 \leq i \leq s$, partition $P_i = \bigcup(P_{i,1}, P_{i,2}, \dots, P_{i,t_i})$ such that for all $1 \leq j \leq t_i - 1$, $k \leq |\bigcup\{H_v : v \in P_{i,j}\}| < 2k$, and $|\bigcup\{H_v : v \in P_{i,t_i}\}| < k$. This is possible since $|H_v| < k$ for all $v \in P_i$. Collapse each set $P_{i,j}$ into a single vertex $p_{i,j}$ in H, whose associated subgraph in G is $H_{p_{i,j}} := \bigcup\{H_v : v \in P_{i,j}\}$. Since the vertices in $P_{i,j}$ have the same neighbourhood, $\{H_v : v \in V(H)\}$ remains a connected partition of G. Let $E_{\text{big}} = \{p_{i,j} : 1 \leq i \leq s, 1 \leq j \leq t_i - 1\}$ and $E_{\text{small}} = \{p_{i,t_i} : 1 \leq i \leq s\}$. Then every vertex in E_{big} is big and every vertex in E_{small} is small.

Suppose that $|E_{\text{small}}| \geq 2^{\ell^2} \cdot |A|$. Let X be the graph obtained from A by adding a clique on $N(v)$ for each vertex $v \in E_{\text{small}}$. Since distinct vertices in E_{small} have distinct neighbourhoods, this process adds at least $|E_{\text{small}}| \geq 2^{\ell^2} \cdot |A|$ cliques. Thus by Lemma 1, a K_ℓ-model of X can be computed in $\mathcal{O}(|G|)$ time. For every edge $x_i x_j$ in this K_ℓ-model that is in X but not in A, we have $x_i, x_j \in N(v)$ for some $v \in E_{\text{small}}$. Since v is not assigned, there is a vertex $u \in D \cup E^*$ assigned to $\{x_i, x_j\}$, and u is adjacent to both x_i and x_j. Since u is not in the K_ℓ-model, we can include u in the connected subgraph of the K_ℓ-model that contains x_i or x_j, and we obtain a K_ℓ-model in $A \cup D \cup E^*$ (in particular, without the edge $x_i x_j$). Now assume that $|E_{\text{small}}| < 2^{\ell^2} \cdot |A|$.

Step 6: We have now partitioned $V(H)$ into sets $A \cup E_{\text{big}}$ of big vertices, and sets $C \cup D \cup E^* \cup E_{\text{small}}$ of small vertices. We have proved that $|C| < |A|$, $|D| + |E^*| < 2^{\ell-3} \cdot |A|$, and $|E_{\text{small}}| < 2^{\ell^2} \cdot |A|$. Thus the number of small vertices is less than $(1 + 2^{\ell-3} + 2^{\ell^2} + 1) \cdot |A| \leq 2^{\ell^2+\ell-2} \cdot |A|$. By definition, the number of big vertices in H is at most $|G| \cdot k^{-1}$. In particular, $|A| \leq |G| \cdot k^{-1}$. Thus $|H| \leq 2^{\ell^2+\ell-1} \cdot |G| \cdot k^{-1}$. Moreover, every $|H_v| < 2k$ for every vertex $v \in V(H)$. □

References

[1] NOGA ALON, PAUL SEYMOUR, AND ROBIN THOMAS. A separator theorem for nonplanar graphs. *J. Amer. Math. Soc.*, 3(4):801–808, 1990.

[2] ERIC D. DEMAINE AND MOHAMMADTAGHI HAJIAGHAYI. Graphs excluding a fixed minor have grids as large as treewidth, with combinatorial and algorithmic applications through bidimensionality. In *Proc. 16th Annual ACM-SIAM Symp. on Discrete Algorithms* (SODA '05). ACM, 2005.

[3] MARTIN GROHE. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003.

[4] ALEXANDR V. KOSTOCHKA. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica*, 4(4):307–316, 1984.

[5] RICHARD J. LIPTON AND ROBERT E. TARJAN. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36(2):177–189, 1979.

[6] WOLFGANG MADER. Homomorphieeigenschaften und mittlere Kantendichte von Graphen. *Math. Ann.*, 174:265–268, 1967.

[7] SERGE PLOTKIN, SATISH RAO, AND WARREN D. SMITH. Shallow excluded minors and improved graph decompositions. In *Proc. 5th Annual ACM-SIAM Symp. on Discrete Algorithms* (SODA '94), pp. 462–470. ACM, 1994.

[8] NEIL ROBERTSON AND PAUL D. SEYMOUR. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B*, 63(1):65–110, 1995.

[9] ARNOLD L. ROSENBERG AND LENWOOD S. HEATH. *Graph separators, with applications*. Frontiers of Computer Science. Kluwer, 2001.

[10] ANDREW THOMASON. An extremal function for contractions of graphs. *Math. Proc. Cambridge Philos. Soc.*, 95(2):261–265, 1984.